
Hilbert Space Embedding for Dirichlet Process Mixtures

Krikamol Muandet

Empirical Inference Department
Max Planck Institute for Intelligent Systems
Tübingen, Germany
krikamol@tuebingen.mpg.de

Abstract

This paper proposes a Hilbert space embedding for Dirichlet Process mixture models via a stick-breaking construction of Sethuraman [6]. Although Bayesian nonparametrics offers a powerful approach to construct a prior that avoids the need to specify the model size/complexity explicitly, an exact inference is often intractable. On the other hand, frequentist approaches such as kernel machines, which suffer from the model selection/comparison problems, often benefit from efficient learning algorithms. This paper discusses the possibility to combine the best of both worlds by using the Dirichlet Process mixture model as a case study.

1 Dirichlet Process mixture models

Much of the real-world data cannot be explained by nice simple probability models. Rather, they often come from heterogeneous sources of unknown properties, which require more complex probability models. Mixture modelling is a popular way of representing such heterogeneity, and also forms a basis for many Bayesian probabilistic models. Unfortunately, a long-standing difficulty in mixture modelling is choosing the number of mixture components, i.e., the number of sources from which the data are generated. Dirichlet Process mixture model (DPMM) allows for the apriori unbounded number of components whose values can be inferred from the observed data.

As a basis of DPMM, we first give a formal definition of the Dirichlet Process (DP), taken from [3].

Definition 1 (Dirichlet Process). *A Dirichlet Process is a distribution of a random probability measure G over a measurable space (Ω, \mathcal{B}) , such that for any finite partition (A_1, \dots, A_r) of Ω (i.e., $\Omega = \coprod_{i=1}^r A_i$, where \coprod means disjoint union and $A_i \in \mathcal{B}$), we have*

$$(G(A_1), \dots, G(A_r)) \sim \text{Dir}(\alpha G_0(A_1), \dots, \alpha G_0(A_r))$$

where $G(A_i) = \int_{A_i} dG$ and $G_0(A_i) = \int_{A_i} dG_0$ for $i = 1, \dots, r$.

Generally speaking, the DP is a distribution over probability measures. Each draw G from a DP can be interpreted as a random distribution, whose sample path is probability measure with probability one. The base distribution G_0 can be thought of as the mean of the DP, whereas the strength parameter α can be regarded as an inverse-variance.

The DP has received much attention and has been extensively studied in the past few years, especially in Bayesian nonparametrics community. Several scenarios have been proposed to show the existence of the DP. For example, Blackwell and MacQueen used the Polya urn scheme to show that the distributions sampled from a DP are discrete almost surely [1]. Equivalent to the extended Polya urn scheme is a Chinese restaurant process (CRP), which is a random process where n customers sit in a Chinese restaurant with an infinite number of tables. Moreover, one may look at draws from a DP as a weighted sum of point masses. This point was made precise by the stick-breaking construction of Sethuraman [6].

In this paper, we resort to this constructive way of forming G . It can be described by the generative process:

$$\beta_i \sim \text{Beta}(1, \alpha), \quad \pi_i = \beta_i \prod_{k=1}^{i-1} (1 - \beta_k), \quad \theta_i \sim G_0, \quad G = \sum_{i=1}^{\infty} \pi_i \delta_{\theta_i}.$$

The following theorem establishes the connection between the stick-breaking construction and the Dirichlet process given in the Definition 1.

Theorem 1. *The stick-breaking construction gives the same probability measure over all random measures on the measurable space (Ω, \mathcal{B}) with the Dirichlet Process with same parameter α and G_0 .*

By mean of the stick-breaking construction, we consider the Dirichlet Process mixture model (DPMM) of the form $\sum_{i=1}^{\infty} \pi_i f_{\theta_i}(x)$, which is a mixture of distributions having the same parametric form f but differing in their parameters. Like many statistical models, exact inference in the DPMM is intractable, and thereby efficient approximate inferences are needed. The most popular inference methods for DPMM are Markov chain Monte Carlo (MCMC), variational Bayesian (VB), and collapsed variational methods. Unlike most previous approaches in nonparametric Bayesian, we study a new approach by employing the Hilbert space embedding. This approach leads to the kernel-based inference for DPMM.

2 Hilbert space embedding for Dirichlet Process mixtures

If we consider the base measure G_0 to be the distribution over the parameter space Θ and let f_{θ} , $\theta \sim G_0$ denote the density function parametrized by θ . Each draw from the DPMM defines the density function $F_{\theta}(x) = \sum_{i=1}^{\infty} \pi_i f_{\theta_i}(x)$. We will represent the probability distribution with density F_{θ} as $\mathbb{P}_{\pi, \theta}$ and represent the set of all $\mathbb{P}_{\pi, \theta}$ by $\mathfrak{P}_{\alpha, \Theta}$.

Let \mathcal{H} be the reproducing kernel Hilbert space (RKHS) with a reproducing kernel k . Assume that $k(x, x)$ is bounded for all x . Then, the Dirichlet Process Mixture Embedding (DPME) is defined as

$$\Upsilon : \mathfrak{P}_{\alpha, \Theta} \longrightarrow \mathcal{H}, \quad \mathbb{P}_{\pi, \theta} \longmapsto \int k(x, \cdot) d\mathbb{P}_{\pi, \theta}(x) \triangleq \sum_{i=1}^{\infty} \pi_i \int k(x, \cdot) df_{\theta_i}(x)$$

We will denote the embedding of $\mathbb{P}_{\pi, \theta}$ by $\Upsilon[\mathbb{P}_{\pi, \theta}]$. Since we have $\sum_{k=1}^{\infty} \pi_k = 1$ almost surely and $k(x, \cdot) < \infty$ for all $x \in \mathcal{X}$, it follows that $\|\Upsilon[\mathbb{P}_{\pi, \theta}]\|_{\mathcal{H}}^2 < \infty$. Therefore, the DPME is well-defined.

Unfortunately, working directly with the DPME is cumbersome because of an infinite sum. Ishwaran and James [4] made an important observation that a truncation of the stick-breaking representation at a sufficiently large T already provides an excellent approximation to the full DPMM model. As a result, we propose the *truncated Dirichlet Process Mixture Embedding* (tDPME):

$$\Upsilon : \mathfrak{P}_{\alpha, \Theta, T} \longrightarrow \mathcal{H}, \quad \mathbb{P}_{\pi, \theta, T} \longmapsto \int k(x, \cdot) d\mathbb{P}_{\pi, \theta, T}(x) \triangleq \sum_{i=1}^T \pi_i \int k(x, \cdot) df_{\theta_i}(x)$$




Figure 1: The hierarchical structure of the truncated DPME. The truncation level T imposes a hierarchical structure on the class of distributions $\mathfrak{P}_{\alpha, \Theta, T}$. As T increases, the set $\mathfrak{P}_{\alpha, \Theta, T}$ enlarges, giving more flexibility to the model.

The $\mathfrak{P}_{\alpha, \Theta, T}$ and $\mathbb{P}_{\pi, \theta, T}$ denote the truncated version of $\mathfrak{P}_{\alpha, \Theta}$ and $\mathbb{P}_{\pi, \theta}$, respectively, where $T > 0$ is a truncation level. The following theorem presents the RKHS version of the almost-sure truncation known in the nonparametric Bayesian literature.

Theorem 2 (Almost-sure truncation). *Let \mathcal{H} be a reproducing kernel Hilbert space (RKHS) with a reproducing kernel k . Assume that $\|k(x, \cdot)\|_{\mathcal{H}}^2 \leq R$ for all x . The following inequality holds:*

$$\|\Upsilon[\mathbb{P}_{\pi, \theta}] - \Upsilon[\mathbb{P}_{\pi, \theta, T}]\|_{\mathcal{H}}^2 \leq C \cdot \exp(-T/\alpha)$$

where C is an arbitrary constant.

Proof. we have

$$\begin{aligned} \|\Upsilon[\mathbb{P}_{\pi, \theta}] - \Upsilon[\mathbb{P}_{\pi, \theta, T}]\|_{\mathcal{H}}^2 &= \left\| \sum_{i=1}^{\infty} \pi_i \int k(x, \cdot) df_{\theta_i}(x) - \sum_{i=1}^T \pi_i \int k(x, \cdot) df_{\theta_i}(x) \right\|_{\mathcal{H}}^2 \\ &= \left\| \sum_{i=T+1}^{\infty} \pi_i \int k(x, \cdot) df_{\theta_i}(x) \right\|_{\mathcal{H}}^2 \\ &\leq \sum_{i=T+1}^{\infty} \pi_i \int \|k(x, \cdot)\|_{\mathcal{H}}^2 df_{\theta_i}(x) \leq \sum_{i=T+1}^{\infty} \pi_i \int R df_{\theta_i}(x) . \end{aligned}$$

We can see that $\int R df_{\theta_i}(x)$ is finite for all i . Thus, letting $\int R df_{\theta_i}(x) < C$ for all i with some constant C yields

$$\sum_{i=T+1}^{\infty} \pi_i \int R df_{\theta_i}(x) \leq \sum_{i=T+1}^{\infty} \pi_i C = C \left(1 - \sum_{i=1}^T \pi_i \right) \approx C \cdot \exp(-T/\alpha) .$$

The last step of the proof uses the fact that $\sum_{i=1}^T \pi_i = \sum_{i=1}^T (\exp(-\Gamma_{i-1}/\alpha) - \exp(-\Gamma_i/\alpha)) = 1 - \exp(-\Gamma_T/\alpha) \approx 1 - \exp(-T/\alpha)$ where $\Gamma_T = E_1 + E_2 + \dots + E_T$ and $E_i \sim \exp(1)$ (cf. [5]). ■

Theorem 2 asserts that the truncated DPME is close in RKHS norm to the true DPME with a sufficiently large truncation level T . Consequently, by working with the truncated DPME instead of the true DPME, we are not losing much information. Moreover, the bound also suggests how to choose T . That is, for the error to be smaller than δ , one must choose T such that $T > \alpha \ln(\delta/C)$. The effect of setting different truncation level T in the DPME can be seen in Figure 1.

2.1 Optimization

Given observation x_1, x_2, \dots, x_m , we would like to find $\mathbb{P}_{\pi, \theta, T}$ that is as close as possible to the underlying distribution \mathbb{P} of the observation. To accomplish this, we employ the usual Hilbert space embedding of \mathbb{P} given by $\mu_{\mathbb{P}} = \mathbb{E}_{x \sim \mathbb{P}}[k(x, \cdot)]$. The empirical estimate of $\mu_{\mathbb{P}}$ can be computed from observation as $\hat{\mu}_X = \frac{1}{m} \sum_{k=1}^m k(x_k, \cdot)$. Then, the optimization problem can be cast as follow:

$$\min_{\pi \in \mathbb{R}^T} \|\hat{\mu}_X - \Upsilon[\mathbb{P}_{\pi, \theta, T}]\|_{\mathcal{H}}^2 \quad \text{subject to} \quad \pi^\top \mathbf{1} = 1, \pi_i \geq 0. \quad (1)$$

To prevent overfitting, we introduce a regularizer $\Omega(\pi) = \frac{1}{2} \|\pi\|^2$ with a regularization constant $\varepsilon > 0$. Substituting $\hat{\mu}_X$ and $\Upsilon[\mathbb{P}_{\pi, \theta, T}]$ back into (1) yields a quadratic programming (QP) for π :

$$\min_{\pi \in \mathbb{R}^T} \frac{1}{2} \pi^\top (\mathbf{S} + \varepsilon \mathbf{I}) \pi - \mathbf{R}^\top \pi \quad \text{subject to} \quad \pi^\top \mathbf{1} = 1, \pi_i \geq 0,$$

where \mathbf{I} is the identity matrix, $\mathbf{S} \in \mathbb{R}^{T \times T}$ and $\mathbf{R} \in \mathbb{R}^T$ are given by $\mathbf{S}_{ij} = \langle \mu[f_{\theta_i}], \mu[f_{\theta_j}] \rangle_{\mathcal{H}}$ and $\mathbf{R}_j = \langle \hat{\mu}_X, \mu[f_{\theta_j}] \rangle_{\mathcal{H}}$, respectively, and $\mu[f_{\theta_i}] = \int k(x, \cdot) d f_{\theta_i}(x)$. Note that our optimization problem is similar to the one in [7]. Thus, due to space constraint, we ask the readers to consult [7] on how to compute \mathbf{S} and \mathbf{R} as well as the detail on how to perform an optimization.

The optimization problem we use here is conceptually similar to the variational methods for DPMM [2]. That is, we are minimizing the distance between the approximate quantity $\mathbb{P}_{\pi, \theta, T}$ and the true quantity \mathbb{P} . Moreover, both MCMC and VB require access to the latent variables associated with observations in order to perform an inference, whereas our approach does not require access to the latent variable whatsoever during an inference. The values of the latent variables, on the other hand, are computed as a postprocessing step.

3 Discussions

We are investigating some open questions related to the proposed kernel-based inference of the DPMM. For example, it is vital to understand how the solution of the above optimization problem relates to the solution of the standard inference methods such as maximum likelihood and MAP of the DPMM. Is there a kernel k for which these solutions coincide? What is the effect of choosing different kernel k ? and what is the connection of our approach to the basic k-mean algorithm? The answers to these questions will be the mutual benefit of researchers in kernel methods and Bayesian nonparametrics.

References

- [1] D. Blackwell and J. B. Macqueen. Ferguson distributions via pólya urn schemes. *The Annals of Statistics*, 1:353–355, 1973.
- [2] D. M. Blei and M. I. Jordan. Variational inference for dirichlet process mixtures. *Bayesian Analysis*, 1:121–144, 2005.
- [3] T. S. Ferguson. A bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):pp. 209–230, 1973.
- [4] H. Ishwaran and James. Gibbs sampling methods for stick breaking priors. *Journal of the American Statistical Association*, pages 161–173, Mar. 2001.
- [5] H. Ishwaran and M. Zarepour. Exact and approximate sum representations for the dirichlet process. *Can J Statistics*, 30(2):269–283, 2002.
- [6] J. Sethuraman. A constructive definition of dirichlet priors. *Statistica Sinica*, 4:639–650, 1994.
- [7] L. Song, X. Zhang, A. Smola, A. Gretton, and B. Schölkopf. Tailoring density estimation via reproducing kernel moment matching. In *Proceedings of the 25th international conference on Machine learning, ICML '08*, pages 992–999, New York, NY, USA, 2008. ACM.